ELSEVIER

European Journal of Operational Research 85 (1995) 132–148

EUROPEAN JOURNAL OF OPERATIONAL RESEARCH

Theory and Methodology

Period and phase of customer replenishment: A new approach to the Strategic Inventory/Routing Problem [1]

Ian R. Webb [a,*], Richard C. Larson [b]

[a] Department of Operations and Information Management, School of Business Administration, University of Connecticut, Storrs, CT 06269, USA

[b] Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA, USA

Received May 1993

Abstract

The delivery of inventory to customers involves the use of vehicles for which purchase or lease agreements may need to be signed months or even years before the start of actual delivery operations. Such long lead times make it impractical to build a fleet incrementally and motivated study of the Strategic Inventory Routing Problem (SIRP). SIRP focuses on estimating, in advance of the start of actual delivery operations, the minimum size (or cost) vehicle fleet required to supply inventory from a central depot to spatially dispersed customers. A fundamental difference between SIRP and the associated tactical IRP's (route an existing fleet, given knowledge of actual customer inventory levels) is that all possible realizations of the tactical problem must be considered, at least implicitly, when solving SIRP. This paper generalizes Larson's approach for SIRP through the use of period and phase of customer replenishment as additional decision variables. Routing solutions based on customer-specific period and phase of replenishment are developed for a simple model of the tactical routing problems the fleet will eventually encounter. Estimates of the fleet size required are developed on the basis of these routing solutions. The period/phase approach can be generalized to take long-term operating costs into account. Computational tests show that the new approach yields significant reductions in solution cost when the vehicle is large enough to replenish several customers in a single trip and/or when there is significant variation in the maximum inter-replenishment intervals of the customers.

Keywords: Strategic Inventory Routing Problem

Inventory Routing Problems (IRP's) arise when both inventory and routing considerations are incorporated in a model to adequately capture the characteristics of a distribution system. Typical routing considerations include constraints on route length and/or duration, vehicle carrying capacity, vehicle/location compatibility and the times at which deliveries may be made. Inventory considerations include storage capacity, consumption characteristics and the consequences of stockouts at the customers. IRP's first appeared in the literature in the 1970's [2,12], however the bulk of the existing

* Corresponding author.

[1] This research was supported in part by NSF grants SES-8411871 and 8724649-ECS.

SSDI 0377-2217(93)E0232-M

research on problems of this type has appeared more recently (e.g. [1,3–11,13,14]). Practical optimization algorithms for realistically sized IRP's have not been found and considerable ingenuity has been needed even to develop good heuristics.

There are two different types of IRP's – the Strategic IRP (SIRP) and the Tactical IRP (TIRP). The strategic version of the problem is motivated by the long lead times (months or even years) between the signing of purchase or lease agreements and the availability of vehicles for delivery operations. These long lead times make it impractical to build a fleet incrementally. SIRP focuses on estimating the minimum size (or cost) vehicle fleet required to supply inventory (e.g. industrial gases, fuel, consumer goods) from a central depot to spatially dispersed customers when only the probability distribution for the per unit time demand at each customer is known. TIRP, the related operational or tactical version of the problem, deals with routing an existing vehicle fleet to supply customers whose actual demands for replenishment are known or can be estimated. A major difference between the strategic and tactical versions of the problem is that in solving SIRP all possible realizations of TIRP must, at least implicitly, be considered simultaneously.

The majority of work on IRP's has focussed on TIRP. An exception is the SIRSA heuristic developed by Larson [11] for a version of SIRP faced by the New York City Department of Environmental Protection in planning the acquisition of a barge fleet. SIRSA has a Clarke-Wright savings framework and estimates the minimum required fleet size by dividing the customers into a set of clusters. A basic assumption in SIRSA is that within any cluster, all replenishments are made on a single route visiting all customers in the cluster. When the solution developed for the New York City problem was examined, it was noticed that the Jamaica Bay location was visited several times a week, even though it needed to be visited only once every forty days. These frequent visits were the result of the inclusion in the cluster of other locations which needed to be visited frequently. Since a single route was used to service the entire cluster, the replenishment intervals were determined by the customer requiring the most frequent replenishments. The work reported here was motivated by this inefficiency and is based on the use of period and phase of individual customer replenishment as additional decision variables. Period and phase of replenishment were first discussed in work on several versions of the Period Vehicle Routing Problem [3,10,12,13]. Introducing these variables enlarges the set of feasible replenishment schemes, so the resulting period/phase heuristic is a generalization of SIRSA.

The paper is organized as follows. Section 1 describes the SIRP model and reviews the SIRSA heuristic developed in [11]. Section 2 presents a simple example motivating the work in this paper and defines the new decision variables – period and phase of customer replenishment. A period/phase heuristic which exploits differences between the consumption and storage characteristics of *individual* customers is developed in Section 3. This heuristic is a generalization of SIRSA and uses the same savings framework. The relationship between our work on SIRP and earlier work on the PVRP is also discussed. Section 4 compares the performance of the period/phase and SIRSA heuristics on a set of sixty test problems. The period/phase approach yields significant reductions in solution cost when the vehicle is large enough to replenish several customers in a single trip and/or when there is significant variation in the maximum inter-replenishment intervals of the customers.

1. Problem statement and previous solution approach

SIRP involves a single depot (with an inexhaustible supply of inventory) serving a set $\mathscr{S}$ of N customers. All point-to-point travel distances d_{ij} are deterministic and known. Customer i has inventory storage capacity C_i and maximum acceptable probability of stockout between successive replenishments p_i. If a customer stocks out, replenishment occurs on the next regularly scheduled visit to the customer. The time to load or unload q units of inventory at location i is given by a holding time function $H_i(q)$.

The daily consumption of inventory at any customer is a random variable with known mean m_i and variance σ_i^2. Demands on different days are assumed to be mutually independent, as are demands at different locations. The vehicle fleet will be composed of identical vehicles having travel speed V and carrying capacity W. The objective is to determine the minimum number of vehicles which should be acquired to service the customers such that

(i) vehicle carrying capacity is not exceeded on any route, and
(ii) the probability of any customer stocking out between successive replenishments is not greater than the customer's maximum acceptable probability of stockout.

The model above was appropriate for the problem faced in New York City. The assumptions used are similar to those in other inventory routing models and will be realistic for many applications.

Previous approach for SIRP

The SIRSA heuristic developed in [11] involves two steps:

Step 1. Replace the original probabilistic demand SIRP with an equivalent deterministic demand problem DSIRP.
Step 2. Solve DSIRP to find a fleet size estimate.

Step 1: Define an equivalent deterministic demand problem DSIRP

For customer i, define TMAX_i (the *maximum inter-replenishment interval* for customer i) to be the maximum possible time between successive replenishments (to full capacity C_i) if the probability of stockout at customer i prior to the next replenishment is not to exceed p_i. As described in [11] the cumulative demand at customer i over several days can be approximated by a normal random variable. Hence TMAX_i must satisfy

$$\Pr\left\{ Z \geq \frac{C_i - (\mathrm{TMAX}_i) m_i}{\sqrt{(\mathrm{TMAX}_i)\sigma_i^2}} \right\} = p_i \tag{1}$$

where Z is a zero mean, unit variance normal random variable. Eq. (1) is readily solved for TMAX_i. An *effective daily consumption rate* (or 'worst-case' daily demand) $\mu_i > m_i$ can then be calculated for each customer using

$$\mu_i = C_i / \mathrm{TMAX}_i. \tag{2}$$

This effective daily consumption rate at customer i is the consumption rate which, if realized each day between successive complete replenishments TMAX_i days apart, would exhaust the inventory of customer i at the instant the next replenishment occurs.

DSIRP, an equivalent deterministic demand version of SIRP, is constructed by having customer i experience their effective daily inventory consumption of μ_i every day. Using the effective daily consumption of each customer is a simple way of enforcing the chance constraints (1) without resorting to a complex mathematical model.

Step 2: Solve the equivalent deterministic demand problem DSIRP

The SIRSA fleet size estimate is based on a set of replenishment routes for DSIRP. These routes should, ideally, be optimal or near-optimal for all possible realizations of the tactical IRP the vehicle fleet may eventually meet. Since the number of such realizations is enormous, it is not practical to use a fully detailed model of the tactical problems in developing these routes. SIRSA generates a set of routes for DSIRP on the basis of a Simplified Tactical Model (referred to here as STM1) which is simple

enough to allow estimation of the fleet size but also retains many fundamental characteristics of the tactical IRP's the vehicle fleet will eventually face. STM1 is based on the following four assumptions:

(A1) The set of customers is divided into permanent disjoint clusters.

(A2) Within a cluster, a *single* minimum length TSP route serves all customers.

(A3) Customers are fully replenished whenever they are visited by a vehicle.

(A4) A route is initiated only when necessary, i.e. when stockout is about to occur at a customer on the route or the accumulated demand for replenishment at the customers on the route is about to exceed vehicle carrying capacity. The second of these conditions forces route initiation because of assumption (A3).

Assumption (A1) is included since any solution to an instance of the tactical IRP effectively divides the customers into clusters. The remaining assumptions define fleet operating policies which are consistent with maximizing the efficiency of vehicle use and estimating the *minimum* fleet size required. Assumption (A2) is an explicit statement of the standard operating policy in VRP's and IRP's. Assumptions (A3) and (A4) are similar to assumptions made in [6].

SIRSA solves DSIRP by developing a set of clusters (and associated routes) with minimum *average vehicle requirement*. For any given cluster Ω, the average vehicle requirement ρ is defined to be

$$\rho = \frac{t}{T} = \frac{\text{Total duration of the route serving } \Omega}{\text{Time between successive initiations of the route serving } \Omega}.$$

Assumptions (A3) and (A4) result in route initiations at regular intervals of

$$T = \min\left\{ \min_{i \in \Omega} \text{TMAX}_i, \quad \frac{W}{\sum_{i \in \Omega} \mu_i} \right\}. \tag{3}$$

The first term in (3) ensures the maximum probability of stockout is not exceeded for any customer, while the second ensures that vehicle carrying capacity is not exceeded. The average vehicle requirement of a cluster can be interpreted as the number of vehicles needed, on average, to satisfy the replenishment requirements of the cluster. For any partition of $\mathcal{S}$ into clusters the vehicle fleet size needed to service $\mathcal{S}$ is estimated by summing the average vehicle requirements of the individual clusters and rounding up the result to the next integer.

The DSIRP clusters and associated fleet estimate are developed using a Clarke and Wright savings approach. The initial solution consists of a separate cluster and route for each customer. At each iteration, the saving (reduction in total average vehicle requirement) associated with each possible pairwise combination of clusters is calculated. The pair of clusters with the largest positive saving are merged. The process terminates when no combination of existing clusters will yield a positive saving.

Adding the vehicle requirements of all clusters assumes that in actual daily operations each vehicle can be scheduled so as to be busy 100% of the time. Although this will not be true in practice, the fleet estimate obtained will be sufficient for actual scheduling problems due to the following sources of excess fleet capacity:

(i) use of effective daily consumption at each customer means that in practice customers will not need to be visited as frequently as indicated in DSIRP;

(ii) the savings algorithm is heuristic and will tend to overestimate fleet size; and

(iii) on any particular day, the scheduler will be able to develop routes based on the current inventory levels of the customers. These routes will be more efficient for that particular day than the STM1 routes on which the fleet estimate is based.

2. Period and phase of customer replenishment for DSIRP

The SIRSA heuristic described above solves DSIRP by adopting simplified tactical model STM1. This section defines new decision variables – period and phase of customer replenishment – and describes a more general Simplified Tactical Model based on these variables.

Motivation

Consider the simple two-customer problem shown in Fig. 1. When STM1 is adopted, the feasible solutions are:
(i) form separate clusters for customers 1 and 2 and use routes with inter-route intervals of, respectively, five and twenty time units, or
(ii) form a single cluster and use the route 0-1-2-0 at intervals of five time units.

These strategies result in the following average vehicle requirements:

$$\rho_{\text{SIRSA (separate clusters)}} = \rho_{\text{cluster 1}} + \rho_{\text{cluster 2}} = \frac{t_{01} + t_{10}}{\text{TMAX}_1} + \frac{t_{02} + t_{20}}{\text{TMAX}_2} = 1.140$$

and

$$\rho_{\text{SIRSA (combined cluster)}} = \frac{t_{01} + t_{12} + t_{20}}{\min\{\text{TMAX}_1, \text{TMAX}_2\}} = \frac{7.7}{5} = 1.540.$$

Now consider the following alternative service strategy: replenish customer 1 on days five, ten and fifteen, and on day twenty replenish both customers 1 and 2. The total travel time for all four routes is $3(t_{01} + t_{10}) + (t_{01} + t_{12} + t_{20})$ while the intervals between each successive pair of route initiations is five days. Repeated over time, this strategy satisfies the customer service and vehicle carrying capacity constraints and results in an average vehicle requirement of

$$\rho_{\text{new}} = \frac{3(t_{01} + t_{10}) + (t_{01} + t_{12} + t_{20})}{4 \cdot \text{TMAX}_1} = \frac{19.1}{20} = 0.955,$$

which is 16.2% lower than that for the best solution possible using STM1. The fleet estimate for the best STM1 solution is two vehicles while for the alternative solution it is one vehicle.

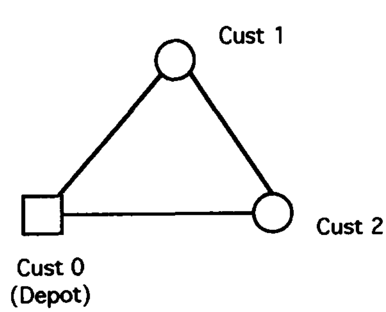

$\text{TMAX}_1 = 5$
$\text{TMAX}_2 = 20$
$W >> C_1 + C_2$

Travel Times $\begin{cases} t_{01} = t_{10} = 1.9 \\ t_{12} = t_{21} = 2.0 \\ t_{20} = t_{02} = 3.8 \end{cases}$

Fig. 1. Two-customer inventory routing problem.

The inefficiency of the STM1 solutions is a result of assumption (A2) and the large difference in the maximum inter-replenishment intervals of the customers. This potential for inefficiency motivated the development of STM2 – a more general Simplified Tactical Model which allows differences between customer inter-replenishment intervals to be exploited.

Period and phase of customer replenishment

Before defining STM2, we introduce the concept of a routeset and define the new decision variables – period and phase of replenishment – used to characterize a routeset.

For any cluster of customers, a *routeset* consists of a number of component routes arranged in a specific order. Each route is a minimum length tour which starts and ends at the depot and visits a subset of the customers in the cluster. The *cardinality* of a routeset is the number of routes it contains. A routeset is repeated over time to service the associated cluster. The alternative solution

$$\mathfrak{R}_1 = \{0\text{-}1\text{-}0,\ 0\text{-}1\text{-}0,\ 0\text{-}1\text{-}0,\ 0\text{-}1\text{-}2\text{-}0\}$$

suggested for the example in Fig. 1 is a routeset of cardinality $|\mathfrak{R}_1| = 4$.

The *period* of customer i (denoted by τ_i) is the *number* of routes after a replenishment of customer i, up to and including the next route on which customer i is replenished. For example, in a cluster consisting of customers 1, 2 and 3 and serviced by the routeset,

$$\mathfrak{R}_2 = \{0\text{-}1\text{-}3\text{-}0,\ 0\text{-}1\text{-}3\text{-}0,\ 0\text{-}1\text{-}2\text{-}3\text{-}0\},$$

customers 1 and 3 have period one (they are serviced on every route) and customer 2 has period three (it is service on every third route).

The *phase* of customer i (denoted by ϕ_i) is the *number* of routes between the beginning of a routeset and the first route replenishing customer i. In the routeset

$$\mathfrak{R}_3 = \{0\text{-}1\text{-}2\text{-}0,\ 0\text{-}1\text{-}3\text{-}0,\ 0\text{-}1\text{-}0,\ 0\text{-}1\text{-}0,\ 0\text{-}1\text{-}2\text{-}3\text{-}0\},$$

customer 1 has a phase of zero and period of one, customer 2 has a phase of zero and period of four and customer 3 has a phase of one and period of three.

Any routeset is completely characterized by its cardinality and the period and phase of each customer it services.

STM2: A more general simplified tactical model

The period/phase algorithm described in Section 3 is based on the following Simplified Tactical Model, which we refer to as STM2. In STM2 each cluster is associated with a routeset rather than a single route. STM2 is defined by the following assumptions:

(A5) The set of customers is divided into permanent disjoint clusters.

(A6) Within a cluster, the customers are serviced by a single routeset. Each customer is visited by a non-empty subset of the component routes.

(A7) Customers are fully replenished whenever they are visited by a vehicle.

(A8) A route is initiated only when necessary, i.e. when stockout is about to occur at a customer on the route or the accumulated demand for replenishment at the customers on the route is about to exceed vehicle carrying capacity.

Assumptions (A5)–(A8) are analogous to assumptions (A1)–(A4) in STM1. Since the single route allowed by (A2) in STM1 is a routeset of cardinality one, the feasible solutions for STM1 are a subset of the feasible solutions for STM2.

3. PPSA: The Period / Phase Savings Algorithm for SIRP

The Period/ Phase Savings Algorithm (PPSA) for SIRP has the same basic structure as SIRSA and involves the following two steps:

Step 1. Replace the original probabilistic demand SIRP with an equivalent deterministic demand problem DSIRP using the equivalent daily demands μ_i.

Step 2. Solve DSIRP (using STM2) to find a period/phase solution with minimum average vehicle requirement.

Step 1 is the same as for SIRSA. Step 2 employs a savings-type heuristic in which the basic computational step involves finding, for any proposed cluster Ω of customers, a routeset with minimum average vehicle requirement. In SIRSA, Step 2 searches only among routesets of cardinality one.

Minimum average vehicle requirement for a cluster

The definition of average vehicle requirement given in Section 1 can now be extended to routesets. Given any cluster Ω and associated routeset $\mathfrak{R}_\Omega$, let t_i be the duration of route i and T_i the interval between the initiations of routes $i-1$ and i. T_1 is the interval between the initiation of the last route in a routeset and the initiation of the first route in the next repetition of the routeset. The *average vehicle requirement* for $\mathfrak{R}_\Omega$ is defined to be

$$\rho_{\mathfrak{R}_\Omega} = \sum_{i=1}^{|\mathfrak{R}|} t_i \Big/ \sum_{i=1}^{|\mathfrak{R}|} T_i \tag{4}$$

where the numerator is the total vehicle time required per repetition of the routeset and the denominator is the time between successive initiations of the routeset. This definition is an extension of the definition in Section 1 since the feasible routes in STM1 are routesets of cardinality one. The average vehicle requirement for a routeset can be interpreted as the number of vehicles required, on average, to implement the routeset.

We now describe procedure **AvVehReqt()**, a two-stage heuristic which, for any specified Ω, searches for a routeset with minimum average vehicle requirement. The first stage determines the cardinality of an initial incumbent routeset in which every route visits all customers in Ω, while the second determines the period and phase of replenishment for each customer.

Stage 1: Determine cardinality of initial incumbent routeset. Given a cluster Ω containing n customers, this step determines the cardinality of the routeset $\mathfrak{R}_\Omega$ serving the cluster. Since the effort required in stage two increases with routeset cardinality, it is desirable to use as few routes as possible. Conversely, since each customer must be replenished at least once in the routeset, using too few routes may force some customers to be replenished more frequently than is necessary (recall the routesets of cardinality one in the example of Fig. 1).

Assumption (A8) ensures that the initiation of a route is a result of either

(i) the accumulated replenishment demand of customers on the route reaching the capacity of the vehicle, or

(ii) the depletion of inventory at one or more of the customers on the route.

From (3), a lower bound on the time between initiations of successive routes in the initial incumbent routeset is

$$\mathrm{LB} = \min\left\{ \min_{i\in\Omega} \mathrm{TMAX}_i,\ \frac{W}{\sum_{i\in\Omega} \mu_i} \right\}. \tag{5}$$

Similarly, an upper bound on the maximum time between successive replenishment of any customer in the cluster is

$$\mathrm{UB} = \min\left\{ \max_{i\in\Omega} \mathrm{TMAX}_i, \quad \max_{i\in\Omega} \frac{W}{\mu_i} \right\}. \tag{6}$$

Since routes are initiated at least LB time units apart and the maximum time between successive replenishments of any customer is UB, the maximum number of routes between successive replenishments (and hence the maximum *period*) of any customer is bounded above by the ratio UB/LB. Employing a routeset of cardinality $|\mathfrak{R}_\Omega| = \lceil \mathrm{UB}/\mathrm{LB} \rceil$ will allow each customer to be replenished at least once in each repetition of the routeset without forcing unnecessarily frequent visits.

Stage 2: Determine period and phase of replenishment for each customer in the cluster. We define an initial incumbent routeset for Ω consisting of $|\mathfrak{R}_\Omega|$ routes, each visiting every customer in Ω, i.e. $\tau_i = 1$ and $\phi_i = 0$ for every customer i. This initial incumbent routeset is modified iteratively. At each iteration the customer/period/phase modification $(i^*, \tau_i^*, \phi_i^*)$ giving maximum reduction in average vehicle requirement is selected by a greedy search over all combinations (i, τ_i, ϕ_i), $\tau_i = 2, \ldots, |\mathfrak{R}|$ and $\phi_i = 0, 1, \ldots, \tau_i - 1$, and for which $\tau_i = 1$ and $\phi_i = 0$ in the current incumbent routeset. The current incumbent routeset is updated by implementing $(i^*, \tau_i^*, \phi_i^*)$. The iterative process continues until no modification reducing the average vehicle requirement is found, or until all customers have had their period and/or phase modified.

Average vehicle requirement of a routeset: The route interval linear program

A routeset is characterized by its cardinality and the period and phase of the customers it serves. Calculating the average vehicle requirement for a routeset involves the sequential calculation of:

(i) the inter-route intervals T_r, $r = 1, 2, \ldots, |\mathfrak{R}_\Omega|$;
(ii) the time intervals between successive replenishments of each customer (a function of the inter-route intervals and the customer periods and phases);
(iii) the delivery quantity to each customer on each route and the total delivery quantity on each route;
(iv) the total duration (travel time plus inventory transfer time) of each route;
(v) the average vehicle requirement (using (4)).

Steps (ii)–(v) are conceptually straightforward, although (iv) involves solving as many as $|\mathfrak{R}_\Omega|$ travelling salesman problems. Calculating the inter-route intervals is not straightforward, as is illustrated by the routeset shown in Fig. 2.

T_1, the interval between route 3 (in the *previous* repetition of the routeset) and route 1 (in the *next* repetition of the routeset) is a function of the accumulated demand at customer 3, which is in turn a function of the intervals T_2 and T_3. Similarly, customer 2 is visited only by route 2, so T_2 is a function of T_1 and T_3. Since T_1 and T_2 are mutually dependent they must be calculated simultaneously. We now show how the inter-route intervals T_r can be found by solving a linear program.

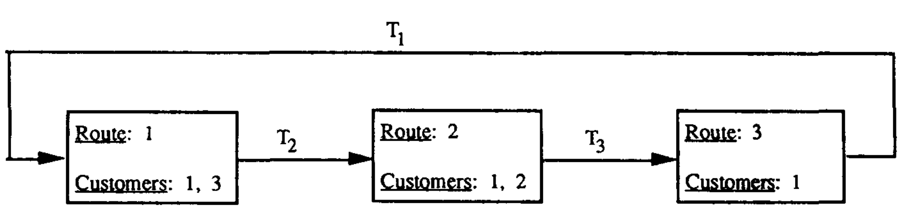

Fig. 2. A simple three-customer, three-route routeset.

The inter-route intervals for the example in Fig. 2 are governed by the set of linear inequalities shown below.

For route 1 (*serving customers* 1 *and* 3):

$T_1 \le \mathrm{TMAX}_1,$ (maximum interval between replenishments of customer 1)

$T_2 + T_3 + T_1 \le \mathrm{TMAX}_3,$ (maximum interval between replenishments of customer 3)

$T_1\mu_1 + (T_2 + T_3 + T_1)\mu_3 \le W.$ (demand on route does not exceed vehicle capacity)

For route 2 (*serving customers* 1 *and* 2):

$T_2 \le \mathrm{TMAX}_1,$ (maximum interval between replenishments of customer 1)

$T_3 + T_1 + T_2 \le \mathrm{TMAX}_2,$ (maximum interval between replenishments of customer 2)

$T_2\mu_1 + (T_3 + T_1 + T_2)\mu_2 \le W.$ (demand on route does not exceed vehicle capacity)

For route 3 (*serving customer* 1):

$T_3 \le \mathrm{TMAX}_1,$ (maximum interval between replenishments of customer 1)

$T_3\mu_1 \le W.$ (demand on route does not exceed vehicle capacity)

Since the goal is efficient vehicle use $\Sigma\, T_r$, the interval between successive repetitions of the routeset, should be as large as possible. The problem of determining the route intervals can therefore be formulated as a Route Interval Linear Program (RILP) in which the T_r are decision variables, the objective is to maximize $\Sigma\, T_r$, and the constraints are the linear inequalities above plus non-negativity of the T_r.

In general, RILP is generated as follows. Let n be the number of customers in Ω, $|\mathfrak{R}_\Omega|$ the number of routes in the routeset and $\Omega_r \subseteq \Omega$ the set of customers served by route r (as determined by $|\mathfrak{R}_\Omega|$ and the customer period and phase combinations being evaluated). For a maximization formulation the decision variables T_r, $r = 1, 2, \dots, |\mathfrak{R}_\Omega|$ each have objective coefficient one. The constraints are of the form $\boldsymbol{AT} \le \boldsymbol{b}$ where $\boldsymbol{T}$ is the vector of decision variables. Constraint matrix $\boldsymbol{A}$ and right-hand side vector $\boldsymbol{b}$ have the structure described below. For route r, $r = 1, 2, \dots, |\mathfrak{R}_\Omega|$:

(a) There is a *customer replenishment interval constraint* for each $i \in \Omega_r$. Denote by (r, i) the row of A containing the constraint associated with route r and customer i. The value of $a_{(r,i),j}$, the coefficient of T_j in this constraint, is

$$a_{(r,i),j} = \begin{cases} 1 & \text{for } j = (L(r, i) + 1), (L(r, i) + 2), \dots r - 1, r, \\ & \text{where } L(r, i) \text{ is the } \textit{last} \text{ route prior to } r \text{ on which } i \text{ was replenished,} \\ 0 & \text{for all other } j, \end{cases}$$

i.e. $a_{(r,i),j}$ is one if and only if customer i is not visited on routes $j, j+1, \dots, r-1$.
The right-hand side value for this constraint is $b_{(r,i)} = \mathrm{TMAX}_i$.

(b) There is a *vehicle capacity constraint*. Denote (r, v) the row of A containing the vehicle capacity constraint for route r. Then $a_{(r,v),j}$, the coefficient of T_j in this constraint, is

$$a_{(r,v)j} = \sum_{k \in M(r,j)} \mu_k \quad \text{where } M(r, j) = \{ i \in \Omega_r \colon a_{(r,i),j} = 1,\ i = 1, 2, \dots, n \}.$$

$M(r, j)$ can be interpreted as the subset of customers in Ω_r who are not replenished on routes j, $j+1, \dots, r-1$. For such customers, demand accumulated over the interval T_j must be met by the delivery on route r.
The right-hand side value for this constraint is: $b_{(r,v)} = W$.

If RILP is infeasible, the set of customer periods and phases being evaluated is also infeasible since it implies at least one negative element of $\boldsymbol{T}$.

RILP involves $|\Re_\Omega|$ decision variables and a maximum of $|\Re_\Omega|\cdot(n+1)$ constraints other than non-negativity constraints on $\boldsymbol{T}$. There are at most n passes through Stage 2 and in each pass $n+1-i$ customers are considered. For customer i, there are at most $|\Re_\Omega|^2$ different (τ_i, ϕ_i) combinations evaluated. Hence the number of instances of RILP solved in any call of **AvVehReqt()** is at most $\frac{1}{2}|\Re_\Omega|^2\cdot(n^2+n)$.

A more formal statement of **AvVehReqt()** can now be given.

Procedure **AvVehReqt(Ω)**

For a given incumbent routeset, let ρ_c be the average vehicle requirement, $U\subseteq\Omega$ the set of customers who have not had their period or phase altered, Δ_ρ the largest possible reduction in ρ_c if the period and/or phase of a single customer in U is altered and $\rho_c(i, \tau_i, \phi_i)$ the average utilization factor if the incumbent routeset is modified by changing the period and phase of customer $i\in U$ to τ_i and ϕ_i respectively. In addition, let Ω_r be the set of customers on route r, $q_{r,j}$ the inventory delivered by route r to customer $j\in\Omega_r$, Q_r the total inventory delivered on route r, $\alpha_{r,j}$ the time since the last replenishment of customer $j\in\Omega_r$, $L(r, j)$ the last route prior to r on which j was replenished and TSP(Ω_r) the length of an optimal TSP tour starting and ending at the depot and visiting all customers in Ω_r.

Stage 1. The cardinality of the initial incumbent routeset serving Ω is $|\Re_\Omega|=\lceil \mathrm{UB}/\mathrm{LB}\rceil$ where UB and LB are calculated using (5) and (6) above. Set $U=\Omega$. For each $i\in\Omega$, set $\tau_i=1$ and $\phi_i=0$. The initial incumbent routeset for Ω has $|\Re_\Omega|$ routes, each of which has $\Omega_r=\Omega$. For $r=1, 2,\dots,|\Re_\Omega|$:

$$T\equiv T_r=\min\left\{\min_{i\in\Omega}\mathrm{TMAX}_i,\quad W/\sum_{i\in\Omega}\mu_i\right\},$$

$$q_i\equiv q_{r,i}=\mu_i\cdot T,\ \forall i\in\Omega,\text{ and } Q\equiv Q_r=\sum_{i\in\Omega}q_i.$$

The average vehicle requirement for this initial incumbent routeset is

$$\rho_c=\left(\mathrm{TSP}(\Omega)/V+H_0(Q)+\sum_{i\in\Omega}H_i(q_i)\right)/T.$$

Stage 2. If $U\neq\emptyset$: Set $\Delta_\rho=0$.

For each $i\in U$:

For $\tau_i=2, 3,\dots,|\Re_\Omega|$ and $\phi_i=0, 1, 2,\dots,\tau_i-1$:

Calculate average vehicle requirement if period and phase of customer i are changed to τ_i and ϕ_i.

Form RILP. Solve for $\boldsymbol{T}$. If RILP is infeasible, go to next $i\in U$.

For $r=1, 2,\dots,|\Re_\Omega|$: Set $Q_r=0$.

For each $j\in\Omega_r$ calculate:

$$\alpha_{r,j}=\sum_{k=L(r,j)+1}^{r}T_k,$$

$$q_{r,j}=\mu_j\cdot\alpha_{r,j},$$

$$Q_r=Q_r+q_{r,j}.$$

The duration of route r is

$$t_r = \mathrm{TSP}(\Omega_r) + H_0(Q_r) + \sum_{k \in \Omega_r} H_k(q_{r,k}).$$

The average vehicle requirement for this routeset is

$$\rho_c(i, \tau_i, \phi_i) = \sum_{r=1}^{|\mathfrak{R}|} t_r \Big/ \sum_{r=1}^{|\mathfrak{R}|} T_r.$$

If $\rho_c - \rho_c(i, \tau_i, \phi_i) > \Delta_\rho$, set $\Delta_\rho = \rho_c - \rho_c(i, \tau_i, \phi_i)$, $i^* = i$, $\tau^* = \tau_i$ and $\phi^* = \phi_i$.

Stage 3. If $\Delta_\rho > 0$, update the incumbent routeset by implementing the period/phase modification (i^*, τ_i^*, ϕ_i^*). Set $\rho_c = \rho_c(i^*, \tau_i^*, \phi_i^*)$ and $U = U \setminus \{i\}$. **Go to Stage 2.**
If $\Delta_\rho \le 0$, **TERMINATE** with the current incumbent routeset. The minimum average vehicle requirement for cluster Ω is ρ_c.

The Period / Phase Algorithm

A formal statement of PPSA can now be given. A collection of routesets will be said to be a *routeset cover* if it services all customers without violating customer stockout or vehicle carrying capacity constraints. Let J be the number of clusters in a routeset cover and c the index for this routeset cover. Define ρ_{cj} to be the average vehicle requirement associated with routeset j in routeset cover c.

Step 1. Define DSIRP. For each customer i, calculate the maximum inter-replenishment interval TMAX_i and effective daily consumption μ_i using Eqs. (1) and (2).

Step 2. Solve DSIRP.

Step 2a. Initialization of Savings Algorithm. The initial routeset cover consists of N routesets. Routeset j consists of a single one-stop route beginning at the depot, visiting customer j, and returning to the depot. Set $J = N$ and $c = 0$. Compute the average vehicle requirement ρ_{cj} for routeset j:
$T(j)$ = Time between successive initiations of routeset $j = \min\{W/\mu_j, \mathrm{TMAX}_j\}$;
$t(j)$ = Total vehicle time required per repetition of routeset j
$= [d_{0j} + d_{j0}]/V + H_0(\mu_j \cdot T(j)) + H_j(\mu_j \cdot T(j))$;
ρ_{0j} = Average vehicle requirement for routeset $j = t(j)/T(j)$.

Step 2b. Calculate savings. Consider pairwise each set of clusters j and k in the current routeset cover. For each pair of clusters, compute:
(i) the minimum average vehicle requirement x_1 associated with retaining the two clusters as separate clusters: $x_1 = \rho_j + \rho_k$;
(ii) the minimum average vehicle requirement x_2 when the two clusters are combined.
Let Ω be the set of customers if clusters j and k are merged. Use procedure **AvVehReqt(Ω)** to calculate x_2;
(iii) Compute the savings possible by combining clusters j and k: $y_{jk} = \max\{x_1 - x_2, 0\}$.

Step 2c. Cluster modification. If at least one of the y_{jk}'s computed in Step 2b is positive, select the largest y_{jk} and revise the current routeset cover, replacing j and k by a single cluster involving the customers from clusters j and k. Eliminate from the previously computed ρ_{ci}'s all those for which $i = j$ or $i = k$. Reindex the routesets in the new routeset cover, increment c by one and decrement J by one. Return to Step 2b.
If no y_{jk} is positive, **TERMINATE**: no further improvement is possible in the current routeset cover. The fleet estimate is $\lceil \sum_{j=1}^{J} \rho_{cj} \rceil$.

Table 1
Sample problem data

Location index	Holding capacity (units)	Effective consumption (units/day)	TMAX (days)	x-coord (km)	y-coord (km)	Setup time (hrs)	Transfer rate (units/hr)
0	∞	1	∞	0.0	0.0	0.750	7000
1	2000	1800	0.900	−181.2	−66.1	0.167	3500
2	380	1690	4.447	−180.5	60.4	0.167	3500
3	680	2085	3.066	−187.9	−11.9	0.167	3500
4	1650	2030	1.230	−143.5	152.7	0.167	3500
5	3570	1515	0.424	−199.3	−14.6	0.167	3500
6	2820	1560	0.553	−36.2	25.3	0.167	3500
7	630	1065	1.690	−174.1	16.4	0.167	3500
8	1285	1880	1.463	−120.6	162.3	0.167	3500
9	580	745	1.284	−149.5	−39.7	0.167	3500

Illustrative example

To illustrate the differences between the solutions generated by SIRSA and PPSA, the nine-customer problem in Table 1 was solved using both heuristics. Vehicle carrying capacity and travel speed are, respectively, 10000 units and 80 km/hr. Inventory transfer times consist of a setup time plus the time to transfer inventory at the indicated rate. The depot is location zero. The SIRSA solution is shown in Table 2. It consists of six clusters each served by a single route. The average vehicle requirement for this solution is 5.06, giving a vehicle fleet estimate of six vehicles. The PPSA solution consists of four clusters. The routeset details are shown in Table 3. The total average vehicle requirement for this solution is 4.73 (a reduction of just over 6.5% cf. the SIRSA figure) and a vehicle fleet estimate of five vehicles (cf. six vehicles for SIRSA).

Minimizing long-term average fleet cost

PPSA has been developed in the context of minimizing the number of vehicles acquired. Minimizing the direct cost of fleet acquisition is a straightforward extension. Since the objective may involve minimizing some combination of acquisition and operating costs, we now indicate how PPSA can be extended to include approximate long-term operating costs per unit time.

For a given cluster an associated routeset is a solution for STM2, which includes all major factors affecting the tactical IRP's except actual customer inventory levels. The best routeset for a cluster can therefore be regarded as a set of 'model' routes which will be close to optimal for most tactical problems eventually encountered in the cluster. Even though the routes used in practice will vary to reflect actual

Table 2
SIRSA solution for sample problem – routeset details

Cluster	Route	Customer sequence	Length (km)	Duration (hours)	Interval since prior route (days)	Avge. vehicle requirement
1	1	0-1-5-0	447.2	7.7	0.424	2.267
2	2	0-2-0	380.8	6.4	4.447	0.180
3	3	0-7-3-9-0	408.8	7.4	1.284	0.721
4	4	0-4-0	419.2	7.0	1.230	0.714
5	5	0-6-0	88.3	2.7	0.553	0.608
6	6	0-8-0	404.8	6.8	1.463	0.579

Table 3
PPSA solution for sample problem – Routeset details

Cluster	Route index	Customer sequence	Length h (km)	Duration (hours)	Interval since prior route (days)	Avge. vehicle requirement
1	1	0-1-5-3-0	447.5	8.7	0.424	
	2	0-5-0	399.7	6.6	0.424	
	3	0-1-5-0	447.2	8.1	0.424	
	4	0-5-0	399.7	6.6	0.424	
	5	0-1-5-0	447.2	8.1	0.424	
	6	0-5-0	399.7	6.6	0.424	
	7	0-1-5-0	447.2	8.1	0.424	2.213
2	1	0-7-9-0	390.8	6.6	1.284	
	2	0-2-7-9-0	450.7	8.0	1.284	0.711
3	1	0-4-0	419.2	7.0	1.230	0.714
4	1	0-6-0	88.3	2.7	0.553	
	2	0-6-0	88.3	2.7	0.553	
	3	0-6-8-0	407.3	7.4	0.357	1.094

customer inventories, the model routes can be used as indicated below to calculate approximations to the long-term operating costs incurred in supplying the customers in the cluster.

For any specified cluster and routeset, the associated vehicle travel distance per unit time can be calculated from TSP(Ω_r) and T_r, the lengths and inter-route intervals for the component routes. Given travel cost data, the long-term travel cost per unit time for the routeset can be calculated and used as an approximation to the long-term travel cost per unit time for servicing the cluster. Similarly, the T_r give the lengths of the intervals between successive replenishments of each customer. Since the cumulative inventory consumption at customer i over several time periods is a normal random variable, the probability of stockout at customer i between any given pair of successive replenishments of i can be calculated. The expected number of stockouts per repetition of the routeset can be calculated for each customer. If a cost per stockout is specified for each customer, the long-term expected stockout cost per unit time can be calculated for the routeset. This can be used as an approximation for the actual long-term stockout cost per unit time for the cluster.

Relationship between current work and the PVRP

Period and phase of replenishment were first discussed in several papers [3,10,12,13] dealing with a vehicle routing problem for which the routing schemes used are similar in character to those generated in PPSA. This problem is referred to as the Period (or Periodic) Vehicle Routing Problem (PVRP) since the length of the planning horizon (or replenishment cycle) is a constant specified in the problem statement.

There are significant differences between the work in the current paper and that in the papers cited above. The first is that the current paper is motivated by a different problem. While the four works cited above focus on the daily routing of an existing vehicle fleet over a specified planning horizon of several days, our work deals with fleet planning months or years in advance of actual daily operations. A second difference is that in [3,10,12,13] period and phase are measured in terms of numbers of time units. In our approach, period and phase are measured in terms of numbers of routes. A final difference is that in the model studied here, period and phase of customer replenishment are decision variables determined solely on the basis of vehicle carrying capacity, customer holding capacity and consumption characteristics, and the relative locations of the customers. In the PVRP papers the length of the planning period is

specified in the problem statement and restrictions are placed on the period and/or phase of replenishment. In [3], the period of replenishment for each customer is specified explicitly as a problem input and only phase of replenishment is a decision variable. In [13] a set of allowable k-day delivery day combinations for each customer is specified and the problem involves choosing one combination for each customer. Russell [12] studies a version of PVRP where the number of times a customer is to be replenished during the planning horizon is specified as an input. A similar restriction is used in [10] where it is assumed that each customer i is replenished exactly once in each interval of T_i days over a planning horizon of m days, where m/T_i is assumed to be integer and m and T_i are specified in the problem statement. These additional restrictions on period and/or phase of replenishment reduce the computation required to find a solution, but may also rule out potentially attractive solutions.

4. Computational results

The New York City data was unavailable. Testing of PPSA was carried out using sixty generated data sets. The test problems were designed to compare the performance of PPSA with that of SIRSA and to identify how any superiority of PPSA is related to two simple characteristics of the problem data:
(i) RV, the ratio of the vehicle carrying capacity to the average customer holding capacity, and
(ii) RT, the ratio of the largest maximum inter-replenishment interval to the smallest maximum inter-replenishment interval among all customers, i.e.

$$\mathrm{RT} = \max_{i=1,2,\dots N}\{\mathrm{TMAX}_i\} \Big/ \min_{i=1,2,\dots N}\{\mathrm{TMAX}_i\}.$$

These two ratios where chosen since:
(i) PPSA is designed to exploit differences between the inter-replenishment intervals of customers. It is reasonable to expect that as these differences increase, PPSA should perform better relative to SIRSA.
(ii) since there is usually a fixed cost associated with each stop a vehicle makes, a large cluster (i.e. one with many customers) is usually inefficient for small vehicles. For larger vehicles, larger clusters may be efficient. In larger clusters, it is more likely that the range of inter-replenishment intervals will be larger, with a corresponding improvement in the performance of PPSA relative to SIRSA.

Two sets of base test problems were used. The first set (Uniform Data sets UD1–UD5) contained five base problems, each involving one hundred customers randomly located over a square distribution region of side 200 km centered on the depot. The second set (Clustered Data sets CD1 through CD5) contained five base problems, each having ten seed points randomly located in the distribution region. A cluster of ten customers randomly located in a square of side 50 km was centered on each seed point. For all base data sets, customer holding capacities were uniformly distributed between 5000 and 35000 gallons. The clustered problems were included since in actual problems customer locations may not be uniformly scattered over the distribution region. An average TMAX value, $\mathrm{TMAX}_{\mathrm{av}}$, was chosen for each base data set. For base data sets one through five the $\mathrm{TMAX}_{\mathrm{av}}$-values chosen were, respectively, 2.5, 3.0, 3.5, 4.0 and 4.5 days.

Each base problem was solved for RV = 4 and 8 and RT = 5.0, 10.0 and 20.0, giving a total of sixty test runs. The TMAX-values of the customers were uniformly distributed in the range $[2 \cdot \mathrm{TMAX}_{\mathrm{av}}/(\mathrm{RT}+1),\ 2 \cdot \mathrm{RT} \cdot \mathrm{TMAX}_{\mathrm{av}}/(\mathrm{RT}+1)]$ to give the required $\mathrm{TMAX}_{\mathrm{av}}$- and RT-values. Vehicle speed was 80 km/hr and vehicle capacity was set equal to the product of RV and the average of the C_i. The holding time functions used were

$$H_0(q) = 0.75 + 1 \times 10^{-5} q \quad \text{for the depot,}$$

$$H_i(q) = 0.167 + 2 \times 10^{-5} q \quad \text{for the customers,}$$

where $H(.)$ is in hours and q in gallons.

Table 4
Computational results

Base data	RT	RV = 4						RV = 8					
		Avge. vehicle requirement		Imp. (%)	Fleet redn.	Time (seconds)		Avge. vehicle requirement		Imp (%)	Fleet redn.	Time (seconds)	
		SIRSA	PPSA			SIRSA	PPSA	SIRSA	PPSA			SIRSA	PPSA
UD1	5.0	8.0737	7.9928	1.0	1	912	10661	6.6879	6.4851	3.0	–	2054	12575
	10.0	9.6804	9.3265	3.7	–	662	18007	8.4389	7.8622	6.8	1	1941	64048
	20.0	11.766	11.608	1.3	–	927	11612	10.846	9.9177	8.6	1	1558	20353
UD2	5.0	14.546	14.415	0.9	–	749	9064	11.710	11.760	0.2	–	1895	6451
	10.0	17.405	17.426	−0.1	–	921	10145	14.967	14.098	5.8	–	1847	38976
	20.0	21.968	20.939	4.7	1	794	21325	19.184	17.930	6.5	2	1271	84633
UD3	5.0	8.5675	8.4512	1.4	–	1049	10885	7.1494	7.0164	1.9	–	3682	5820
	10.0	10.760	10.402	3.3	–	1378	15035	9.5424	9.0072	5.6	–	2049	29954
	20.0	14.273	13.513	5.3	1	1395	16886	13.358	12.042	9.9	1	2179	52851
UD4	5.0	12.440	12.267	1.4	–	1164	8066	10.266	10.217	0.5	–	3228	6188
	10.0	15.394	15.047	2.3	–	831	11060	13.588	12.839	5.5	1	2159	19518
	20.0	18.217	17.738	2.6	1	1079	10892	16.850	15.246	9.5	1	2243	45271
UD5	5.0	9.5080	9.4243	0.9	–	970	9985	7.8272	7.6807	1.9	–	2528	6968
	10.0	11.317	10.915	3.5	1	1110	14314	9.8594	9.2112	6.6	–	1768	19741
	20.0	13.915	13.109	5.8	–	948	19042	12.619	11.544	8.5	1	1624	31052
CD1	5.0	7.4731	7.3935	1.1	–	769	23221	6.0842	5.9771	1.8	1	4269	15720
	10.0	9.0214	8.8316	2.1	1	865	14207	7.7891	7.4365	4.5	–	1560	26501
	20.0	11.128	10.951	1.6	1	1240	11762	9.8536	9.4912	3.7	–	1525	60569
CD2	5.0	14.209	13.967	1.7	1	645	16601	10.952	10.632	2.9	–	1318	17007
	10.0	17.260	16.943	1.8	1	844	28099	14.166	13.206	6.8	1	2247	78312
	20.0	21.755	20.651	5.2	1	1175	52247	19.848	18.098	8.8	1	2073	78964
CD3	5.0	7.6714	7.6462	0.3	–	1005	17720	6.2817	6.1317	2.4	–	1362	11508
	10.0	10.070	9.8189	2.5	1	1208	25902	8.0517	7.9925	0.7	1	2395	34771
	20.0	13.328	12.983	2.6	1	1527	22451	11.625	10.475	9.9	1	1986	98265
CD4	5.0	11.693	11.471	1.9	–	779	20649	9.1786	9.2294	−0.6	–	1633	7274
	10.0	14.377	14.019	2.5	–	799	15180	11.275	11.085	1.7	–	1838	14061
	20.0	17.405	17.082	1.9	–	887	14282	14.555	13.552	6.9	1	1796	42886
CD5	5.0	8.6079	8.4893	1.4	–	969	14670	7.0140	6.7995	3.1	1	2089	10400
	10.0	10.600	10.483	1.1	–	784	16324	8.5472	8.1387	4.8	–	2443	30467
	20.0	12.655	12.303	2.8	–	1145	17397	10.996	10.104	8.1	–	1880	52721

The results of the tests are shown in Table 4 and summarized in Table 5. The primary criterion on which solutions were compared is average vehicle requirement (which will usually be non-integer) rather than fleet size estimate (which is always integer). Average vehicle requirement is a better criterion for *comparison* since it is insensitive to the size of the fleets involved while fleet size (obtained by rounding up the average vehicle requirement) is not. For example, if the SIRSA average vehicle requirement is

Table 5
Average PPSA reduction in average vehicle requirement

Problem type	RV = 4			RV = 8		
	RT = 5	RT = 10	RT = 20	RT = 5	RT = 10	RT = 20
U	1.1	1.9	3.9	2.0	6.1	8.6
C	1.0	2.0	2.9	2.2	4.1	7.5

99.5 even a small percentage reduction in average vehicle requirement will reduce the resulting fleet estimate by at least one vehicle. If however the SIRSA average vehicle requirement is 9.95 the fleet estimate will not change unless the PPSA average vehicle requirement is at least 9.5% lower than that for SIRSA. Computational times shown are in seconds on a Macintosh IIci desktop computer.

PPSA produced reductions in average vehicle requirement in fifty-eight of sixty test problems. Reductions in fleet size were observed in twenty-six of sixty problems. As expected the performance of PPSA relative to SIRSA improves significantly as the variability in the maximum inter-replenishment intervals of the customers increases and/or vehicle capacity increases relative to average customer holding capacity. When RT = 20 and RV = 8 the reduction in average vehicle requirement averages 8.0% and fleet size reductions occur in eight of ten problems with these characteristics. When RT = 5 and RV = 4 there is a 1% reduction in average vehicle requirement and fleet size is reduced in two of ten problems.

It is worth noting that pursuing reductions in fleet size through the period/phase approach is worthwhile even in problems for which RT and/or RV are low. Reducing a fleet by even a single vehicle can result in substantial capital and operating cost savings (e.g. the vehicles involved in [11] were barges costing approximately $5 million dollars each).

PPSA performed somewhat better on the uniform problems than on the clustered problems, particularly for higher values of RT and RV. When the customers are clustered, the distance between the depot and the cluster will typically be much greater than the distances between customers within the cluster. If frequent visits are made to a customer with a small TMAX value, the additional cost of visiting other customers in the cluster on each visit is relatively small, even if these other customers do not require frequent replenishment. For uniformly distributed customers, customer-to-customer distances are larger relative to depot-to-customer distances and this additional cost is larger. As a result the 'one-route-per-cluster' assumption of SIRSA is not as costly relative to PPSA when customer locations are clustered.

The computational effort required by PPSA was considerably greater than for SIRSA. This is not a serious disadvantage since the potential benefits in term of reduced acquisition and operating costs for the fleet are much greater than the cost of using a desktop computer for the length of time required for PPSA. All but a handful of the test problems could have been solved overnight on a machine which is slow even by desktop computer standards.

References

[1] Bell, W., Dalberto, L., Fisher, M., Greenfield, A., Jaikumar, P., Kedia, P., Macj, R., and Prutzman, P., "Improving the distribution of industrial gases with an on-line computerized routing and scheduling optimizer", *Interfaces* 13 (1983) 4–23.

[2] Beltrami, E., and Bodin, L., "Networks and vehicle routing for municipal waste collection", *Networks* 4 (1974) 65–94.

[3] Christofides, N., and Beasley, J., "The period routing problem", *Networks* 14 (1984) 237–256.

[4] Dror, M., "Modeling vehicle routing with uncertain demands as a stochastic program: Properties of the corresponding solution", *European Journal of Operational Research* 64 (1993) 432–441.

[5] Dror, M., and Ball, M., "A computational comparison of algorithms for the Inventory Routing Problem", *Annals of Operations Research* 4 (1986) 3–23.

[6] Dror, M., and Ball, M., "Inventory/routing: Reduction from an annual to a short-period problem", *Naval Research Logistics* 34/6 (1987) 891–905.

[7] Dror, M., Laporte, G., and Trudeau, P., "Vehicle routing with stochastic demands: Properties and solution frameworks", *Transportation Science* 23/3 (1989) 166–176.

[8] Dror, M., and Levy, L., "A vehicle routing improvement algorithm comparison of a 'greedy' and a matching implementation for inventory routing", *Computers Operations Research* 13/1 (1986) 33–45.

[9] Federgruen, A., and Zipkin, P., "A combined vehicle routing and inventory allocation problem", *Operation Research* 32/5 (1984) 1017–037.

[10] Gaudioso, M., and Paletta, G., "A heuristic for the periodic vehicle routing problem", *Transportation Science* 26/2 (1992) 86–92.

[11] Larson, R., "Transporting sludge to the 106-mile site: An inventory/routing model for fleet sizing and logistics system design", *Transportation Science* 22/3 (1988) 186–198.

[12] Russell, R., and Igo, W., "An assignment routing problem", *Networks* 9 (1979) 1–17.

[13] Tan, C., and Beasley, J., "A heuristic algorithm for the period vehicle routing problem", *Omega* 12 (1984) 497–504.

[14] Trudeau, P., and Dror, M., "Stochastic inventory routing: Route design with stockouts and route failures", *Transportation Science* 26/3 (1992) 171–184.